

Shallow Packings in Geometry ^{*}

Esther Ezra [†]

December 18, 2014

Abstract

We refine the bound on the packing number, originally shown by Haussler, for shallow geometric set systems. Specifically, let \mathcal{V} be a finite set system defined over an n -point set X ; we view \mathcal{V} as a set of indicator vectors over the n -dimensional unit cube. A δ -separated set of \mathcal{V} is a subcollection \mathcal{W} , s.t. the Hamming distance between each pair $\mathbf{u}, \mathbf{v} \in \mathcal{W}$ is greater than δ , where $\delta > 0$ is an integer parameter. The δ -packing number is then defined as the cardinality of the largest δ -separated subcollection of \mathcal{V} . Haussler showed an asymptotically tight bound of $\Theta((n/\delta)^d)$ on the δ -packing number if \mathcal{V} has VC-dimension (or *primal shatter dimension*) d . We refine this bound for the scenario where, for any subset, $X' \subseteq X$ of size $m \leq n$ and for any parameter $1 \leq k \leq m$, the number of vectors of length at most k in the restriction of \mathcal{V} to X' is only $O(m^{d_1} k^{d-d_1})$, for a fixed integer $d > 0$ and a real parameter $1 \leq d_1 \leq d$ (this generalizes the standard notion of *bounded primal shatter dimension* when $d_1 = d$). In this case when \mathcal{V} is “ k -shallow” (all vector lengths are at most k), we show that its δ -packing number is $O(n^{d_1} k^{d-d_1} / \delta^d)$, matching Haussler’s bound for the special cases where $d_1 = d$ or $k = n$. As an immediate consequence we conclude that set systems of halfspaces, balls, and parallel slabs defined over n points in d -space admit better packing numbers when k is smaller than n . Last but not least, we describe applications to (i) spanning trees of low total crossing number, and (ii) geometric discrepancy, based on previous work by the author.

1 Introduction

Let \mathcal{V} be a set system defined over an n -point set X . We follow the notation in [19], and view \mathcal{V} as a set of indicator vectors in \mathbb{R}^n , that is, $\mathcal{V} \subseteq \{0, 1\}^n$. Given a subsequence of indices (coordinates) $I = (i_1, \dots, i_k)$, $1 \leq i_j \leq n$, $k \leq n$, the *projection* $\mathcal{V}_{|I}$ of \mathcal{V} onto I (also referred to as the *restriction* of \mathcal{V} to I) is defined as

$$\mathcal{V}_{|I} = \{(\mathbf{v}_{i_1}, \dots, \mathbf{v}_{i_k}) \mid \mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_n) \in \mathcal{V}\}.$$

With a slight abuse of notation we write $I \subseteq [n]$ to state the fact that I is a subsequence of indices as above. We now recall the definition of the primal shatter function of \mathcal{V} :

Definition 1.1 (Primal Shatter Function [21, 29]). *The primal shatter function of $\mathcal{V} \subseteq \{0, 1\}^n$ is a function, denoted by $\pi_{\mathcal{V}}$, whose value at m is defined by $\pi_{\mathcal{V}}(m) = \max_{I \subseteq [n], |I|=m} |\mathcal{V}_{|I}|$. In other words, $\pi_{\mathcal{V}}(m)$ is the maximum possible number of distinct vectors of \mathcal{V} when projected onto a subsequence of m indices.*

^{*}Work on this paper has been supported by NSF under grant CCF-12-16689 and CCF-11-17336.

[†]Department of Computer Science and Engineering, Polytechnic Institute of NYU, Brooklyn, NY 11201-3840, USA; and School of Mathematics, Georgia Institute of Technology, Atlanta, Georgia 30332, USA; esther@courant.nyu.edu

From now on we say that $\mathcal{V} \subseteq \{0, 1\}^n$ has *primal shatter dimension* d if $\pi_{\mathcal{V}}(m) \leq Cm^d$, for all $m \leq n$, where $d > 1$ and $C > 0$ are constants. A notion closely related to the primal shatter dimension is that of the *VC-dimension*:

Definition 1.2 (VC-dimension [19, 36]). *An index sequence $I = (i_1, \dots, i_k)$ is shattered by \mathcal{V} if $\mathcal{V}|_I = \{0, 1\}^k$. The VC-dimension of \mathcal{V} , denoted by d_0 is the size of the longest sequence I shattered by \mathcal{V} . That is, $d_0 = \max\{k \mid \exists I = (i_1, i_2, \dots, i_k), 1 \leq i_j \leq n, \text{ with } \mathcal{V}|_I = \{0, 1\}^k\}$.*

The notions of primal shatter dimension and VC-dimension are interrelated. By the Sauer-Shelah Lemma (see [33, 35] and the discussion below) the VC-dimension of a set system \mathcal{V} always bounds its primal shatter dimension, that is, $d \leq d_0$. On the other hand, when the primal shatter dimension is bounded by d , the VC-dimension d_0 does not exceed $O(d \log d)$ (which is straightforward by definition, see, e.g., [15]).

A typical family of set systems that arise in geometry with bounded primal shatter (resp., VC-) dimension consists of set systems defined over points in some low-dimensional space \mathbb{R}^d , where \mathcal{V} represents a collection of certain simply-shaped regions, e.g., halfspaces, balls, or simplices in \mathbb{R}^d . In such cases, the primal shatter (and VC-) dimension is a function of d ; see, e.g., [15] for more details. When we flip the roles of points and regions, we obtain the so-called *dual set systems* (where we refer to the former as *primal set systems*). In this case, the ground set is a collection \mathcal{S} of algebraic surfaces in \mathbb{R}^d , and \mathcal{V} corresponds to faces of all dimensions in the *arrangement* $\mathcal{A}(\mathcal{S})$ of \mathcal{S} , that is, this is the decomposition of \mathbb{R}^d into connected open *cells* of dimensions $0, 1, \dots, d$ induced by \mathcal{S} . Each cell is a maximal connected region that is contained in the intersection of a fixed number of the surfaces and avoids all other surfaces; in particular, the 0-dimensional cells of $\mathcal{A}(\mathcal{S})$ are called “vertices”, and d -dimensional cells are simply referred to as “cells”; see [34] for more details. The distinction between primal and dual set systems in geometry is essential, and set systems of both kinds appear in numerous geometric applications, see, once again [15] and the references therein.

δ -packing. The *length* $\|\mathbf{v}\|$ of a vector $\mathbf{v} \in \mathcal{V}$ under the L^1 norm is defined as $\sum_{i=1}^n |\mathbf{v}_i|$, where \mathbf{v}_i is the i th coordinate of \mathbf{v} , $i = 1, \dots, n$. The *distance* $\rho(\mathbf{u}, \mathbf{v})$ between a pair of vectors $\mathbf{u}, \mathbf{v} \in \mathcal{V}$ is defined as the L^1 norm of the difference $\mathbf{u} - \mathbf{v}$, that is, $\rho(\mathbf{u}, \mathbf{v}) = \sum_{i=1}^n |\mathbf{u}_i - \mathbf{v}_i|$. In other words, it is the *symmetric difference distance* between the corresponding sets represented by \mathbf{u}, \mathbf{v} .

Let $\delta > 0$ be an integer parameter. We say that a subset of vectors $\mathcal{W} \subseteq \{0, 1\}^n$ is δ -*separated* if for each pair $\mathbf{u}, \mathbf{v} \in \mathcal{W}$, $\rho(\mathbf{u}, \mathbf{v}) > \delta$. The δ -*packing number* for \mathcal{V} , denote it by $\mathcal{M}(\delta, \mathcal{V})$, is then defined as the cardinality of the largest δ -separated subset $\mathcal{W} \subseteq \mathcal{V}$. A key property, originally shown by Haussler [19] (see also [8, 9, 11, 29, 37]), is that set systems of bounded primal shatter dimension admit small δ -packing numbers. That is:

Theorem 1.3 (Packing Lemma [19, 29]). *Let $\mathcal{V} \subseteq \{0, 1\}^n$ be a set of indicator vectors of primal shatter dimension d , and let $1 \leq \delta \leq n$ be an integer parameter. Then $\mathcal{M}(\delta, \mathcal{V}) = O((n/\delta)^d)$, where the constant of proportionality depends on d .*

We note that in the original formulation in [19] the assumption is that the set system has a finite VC-dimension. However, its formulation in [29], which is based on a simplification of the analysis of Haussler by Chazelle [8], relies on the assumption that the primal shatter dimension is d , which is the actual bound that we state in Theorem 1.3. We also comment that a closer inspection of the analysis in [19] shows that this assumption can be replaced with that of having bounded primal shatter dimension (independent of the analysis in [8]). We describe these considerations in Section 2.

Previous work. In his seminal work, Dudley [11] presented the first application of *chaining*, a proof technique due to Kolmogorov, to empirical process theory, where he showed the bound $O((n/\delta)^{d_0} \log^{d_0}(n/\delta))$ on $\mathcal{M}(\delta, \mathcal{V})$, with a constant of proportionality depending on the VC-dimension d_0 (see also previous work by Haussler [18] and Pollard [32] for an alternative proof and a specification of the constant of proportionality). This bound was later improved by Haussler [19], who showed $\mathcal{M}(\delta, \mathcal{V}) \leq e(d_0 + 1) \left(\frac{2en}{\delta}\right)^{d_0}$ (see also Theorem 1.3), and presented a matching lower bound, which leaves only a constant factor gap, which depends exponentially in d_0 . In fact, the aforementioned bounds are more general, and can also be applied to classes of real-valued functions of finite “pseudo-dimension” (the special case of set systems corresponds to Boolean functions), see, e.g., [18], however, we do not discuss this generalization in this paper and focus merely on set systems \mathcal{V} of finite primal shatter (resp., VC-) dimension.

The bound of Haussler [19] (Theorem 1.3) is in fact a generalization of the so-called Sauer-Shelah Lemma [33, 35], asserting that $|\mathcal{V}| \leq (en/d_0)^{d_0}$, where e is the base of the natural logarithm, and thus this bound is $O(n^{d_0})$. Indeed, when $\delta = 1$, the corresponding δ -separated set should include all vectors in \mathcal{V} , and then the bound of Haussler [19] becomes $O(n^{d_0})$, matching the Sauer-Shelah bound up to a constant factor that depends on d_0 .

There have been several studies extending Haussler’s bound or improving it in some special scenarios. We name only a few of them. Gottlieb *et al.* [14] presented a sharpening of this bound when δ is relatively large, i.e., δ is close to $n/2$, in which case the vectors are “nearly orthogonal”. They also presented a tighter lower bound, which considerably simplifies the analysis of Bshouty *et al.* [6], who achieved the same tightening.

A major application of packing is in obtaining improved bounds on the *sample complexity* in machine learning. This was studied by Li *et al.* [23] (see also [18]), who presented an asymptotically tight bound on the sample complexity, in order to guarantee a small “relative error.” This problem has been revisited by Har-Peled and Sharir [17] in the context of geometric set systems, where they referred to a sample of the above kind as a “relative approximation” (discussed in Appendix C), and showed how to integrate it into an *approximate range counting* machinery, which is a central application in computational geometry. The packing number has also been used by Welzl [37] in order to construct spanning trees of low crossing number (see also [29]) and by Matoušek [28, 29] in order to obtain asymptotically tight bounds in geometric discrepancy. We discuss these applications in the context of the problem studied in this paper in Section 4.

Our result. In the sequel, we refine the bound in the Packing Lemma (Theorem 1.3) so that it becomes sensitive to the length of the vectors $\mathbf{v} \in \mathcal{V}$, based on an appropriate refinement of the underlying primal shatter function. This refinement has several geometric realizations. Our ultimate goal is to show that when the set system is “shallow” (that is, the underlying vectors are short), the packing number becomes much smaller than the bound in Theorem 1.3.

Nevertheless, we cannot always enforce such an improvement, as in some settings the worst-case asymptotic bound on the packing number is $\Omega((n/\delta)^d)$ even when the set system is shallow. We demonstrate such a scenario by considering dual set systems of axis-parallel rectangles and points in the plane, where one can have a large subcollection \mathcal{F} that is both δ -separated and δ -shallow. In this case $|\mathcal{F}| = \Omega((n/\delta)^2)$, which is not any better than the “standard” bound (stated in Theorem 1.3) obtained without the shallowness assumption. See Figure 1 and Appendix A, where we give a more detailed description of this construction to the non-expert reader.

Therefore, in order to obtain an improvement on the packing number of shallow set systems, we may need further assumptions on the primal shatter function. Such assumptions stem from the random sampling technique of Clarkson and Shor [10], which we define as follows. Let \mathcal{V} be our set system. We assume that for any sequence I of $m \leq n$ indices, and for any parameter $1 \leq k \leq m$, the number of vectors in $\mathcal{V}|_I$ of length at most k is only $O(m^{d_1} k^{d-d_1})$, where d is the primal shatter dimension and $1 \leq d_1 \leq d$ is a real parameter.¹ When $k = m$ we obtain $O(m^d)$ vectors in total, in accordance with the assumption that the primal shatter dimension is d , but the above bound is also sensitive to the length of the vectors as long as $d_1 < d$. From now on, we say that a primal shatter function of this kind has the (d, d_1) *Clarkson-Shor property*.

Let us now denote by $\mathcal{M}(\delta, k, \mathcal{V})$ the δ -packing number of \mathcal{V} , where the vector length of each element in \mathcal{V} is at most k , for some integer parameter $1 \leq k \leq n$. By these assumptions, we can assume, without loss of generality, that $k \geq \delta/2$, as otherwise the distance between any two elements in \mathcal{V} must be strictly less than δ , in which case the packing is empty. We also assume that $\delta \leq n/2^{(d_0+1)}$ (where d_0 is the VC-dim), as otherwise the bound on the packing number is a constant that depends on d and d_0 by the Packing Lemma (Theorem 1.3). The choice of this threshold is justified in Section 3 where we present the analysis and show our main result, which we state below:

Theorem 1.4 (Shallow Packing Lemma). *Let $\mathcal{V} \subseteq \{0, 1\}^n$ be a set of indicator vectors, whose primal shatter function has a (d, d_1) Clarkson-Shor property, and whose VC-dim is d_0 . Let δ be an integer parameter between 1 and $n/2^{(d_0+1)}$, k an integer parameter between 1 and n , and suppose that $k \geq \delta/2$. Then:*

$$\mathcal{M}(\delta, k, \mathcal{V}) = O\left(\frac{n^{d_1} k^{d-d_1}}{\delta^d}\right),$$

where the constant of proportionality depends on d .

This problem has initially been addressed by the author in [13] as a major tool to obtain size-sensitive discrepancy bounds in set systems of this kind, where it has been shown $\mathcal{M}(\delta, k, \mathcal{V}) = O\left(\frac{n^{d_1} k^{d-d_1} \log^d(n/\delta)}{\delta^d}\right)$. The analysis in [13] is a refinement over the technique of Dudley [11] combined with the existence of small-size *relative approximations* (see [13] and Appendix C for more details). In the current analysis we completely remove the extra $\log^d(n/\delta)$ factor appearing in the previous bound. In particular, when $d_1 = d$ (where we just have the original assumption on the primal shatter function) or $k = n$ (in which case each vector in \mathcal{V} has an arbitrary length), our bound matches the tight bound of Haussler, and thus appears as a generalization of the Packing Lemma (when replacing VC-dimension by primal shatter dimension).

Theorem 1.4 implies smaller packing numbers for several natural geometric set systems under the shallowness assumption. These set systems are described in detail in Section 4.1.

Next, in Section 4.2 we present an application of Theorem 1.4 to “spanning trees with low total conflict number”, which is based on the machinery of Welzl [37] to construct spanning trees of low crossing number (see also [29]). Here the tree spans \mathcal{V} (representing, say, a set of regions defined over n points in d -space), and the “conflict number” of an edge (u, v) is the symmetric difference distance between u and v . Based on this structure we introduce a general framework to efficiently compute various measures arising in geometric optimization (e.g., diameter, width, radius of the smallest enclosing ball, volume of the minimum bounding box, etc.) in each region represented by \mathcal{V} , where the key idea is to keep the overall number of updates small

¹We ignore the cases where $d_1 < 1$, as it does not seem to appear in natural set systems—see below.

(given a spanning tree of the above kind). In Section 4.3 we show an application in geometric discrepancy, where the goal is to obtain discrepancy bounds that are sensitive to the length of the vectors in \mathcal{V} . Due to the bound in Theorem 1.4 we obtain an improvement over the one presented in [13].

In Section 5 we discuss the geometric interpretation of Theorem 1.4 to dual set systems. In particular, we draw the connection between shallow packings and *shallow cuttings* [27].

Beyond the geometric applications, this paper is primarily an extension of Haussler’s technique [19] to shallow set systems. We note that whereas the analysis of Dudley [11] is fairly simple and intuitive, the analysis of Haussler [19] is much more intricate, and thus the initial effort in this study was to understand Haussler’s analysis, whose conclusions are summarized in Appendix B. We are also aware of the simplification to Haussler’s proof by Chazelle [8], nevertheless, we had to use the observations made in [19] in order to proceed with our analysis. Our main conclusion about the analysis in [19] is given in Inequality (1), which implies that the cardinality of a δ -separated set \mathcal{V} is bounded (up to a factor of $(d_0 + 1)$) by the expected number of vectors in the projection of \mathcal{V} onto a random sample of $O(d_0 n / \delta)$ indices, where d_0 is the VC-dimension. Although this simple observation is not explicitly stated in [19], this relation is a key property used in our analysis.

2 Preliminaries

Overview of Haussler’s Approach. For the sake of completeness, we repeat some of the details in the analysis of Haussler [19] and use similar notation for ease of presentation.

Let $\mathcal{V} \subseteq \{0, 1\}^n$ be a collection of indicator vectors of bounded primal shatter dimension d , and denote its VC-dimension by d_0 . By the discussion above, $d_0 = O(d \log d)$. From now on we assume that \mathcal{V} is δ -separated, and thus a bound on $|\mathcal{V}|$ is also a bound on the packing number of \mathcal{V} . The analysis in [19] exploits the method of “conditional variance” in order to conclude

$$|\mathcal{V}| \leq (d_0 + 1) \mathbf{Exp}_I [|\mathcal{V}_I|] = O(d \log d \mathbf{Exp}_I [|\mathcal{V}_I|]), \quad (1)$$

where $\mathbf{Exp}_I [|\mathcal{V}_I|]$ is the expected size of \mathcal{V} when projected onto a subset $I = \{i_1, \dots, i_{m-1}\}$ of $m - 1$ indices chosen uniformly at random without replacements from $[n]$, and

$$m := \left\lceil \frac{(2d_0 + 2)(n + 1)}{\delta + 2d_0 + 2} \right\rceil = O\left(\frac{d_0 n}{\delta}\right) = O\left(\frac{nd \log d}{\delta}\right). \quad (2)$$

We justify this choice in Appendix B, as well as the facts that $m \leq n$ and I consists of precisely $m - 1$ indices.

For the sake of completeness, we review Haussler approach in Appendix B, and also emphasize some of the properties there, which are fundamental in our view. Moreover, we refine Haussler’s analysis to include two natural extensions: (i) *Obtain a refined bound on $\mathbf{Exp}_I [|\mathcal{V}_I|]$* : This extension is a direct consequence of Inequality (1). In the analysis of Haussler $\mathbf{Exp}_I [|\mathcal{V}_I|]$ is replaced by its upper bound $O(m^d)$, resulting from the fact that the primal shatter dimension of \mathcal{V} (and thus of \mathcal{V}_I) is d , from which we obtain that *for any choice of I , $|\mathcal{V}_I| = O((m - 1)^d) = O(m^d)$* , with a constant of proportionality that depends on d , and thus the packing number is $O((n/\delta)^d)$, as asserted in Theorem 1.3.² However, in our analysis we would

²We note, however, that the original analysis of Haussler [19] does not rely on the primal shatter dimension, and the bound on $\mathbf{Exp}_I [|\mathcal{V}_I|]$ is just $O(m^{d_0})$ due to the Sauer-Shelah Lemma.

like to have a more subtle bound on the actual expected value of $|\mathcal{V}_{|_I}|$. In fact, the scenario imposed by our assumptions on the set system eventually yields a much smaller bound on the expectation of $|\mathcal{V}_{|_I}|$, and thus on $|\mathcal{V}|$. We review this in more detail below. (ii) *Relaxing the bound on m .* We show that Inequality (1) is still applicable when the sample I is slightly larger than the bound in (2), as a stand alone relation, this may result in a suboptimal bound on $|\mathcal{V}|$, however, this property will assist us to obtain local improvements over the bound on $|\mathcal{V}|$, eventually yielding the bound in Theorem 1.4. Specifically, in our analysis, described in Section 3, we proceed in iterations, where at the first iteration we obtain a preliminary bound on $|\mathcal{V}|$ (Corollary 3.2), and then, at each subsequent iteration $j > 1$, we draw a sample I_j of $m_j - 1$ indices where

$$m_j := m \log^{(j)}(n/\delta) = O\left(\frac{d_0 n \log^{(j)}(n/\delta)}{\delta}\right), \quad (3)$$

m is our choice in (2), and $\log^{(j)}(\cdot)$ is the j th iterated logarithm function. Then, by a straightforward generalization of Haussler’s analysis (described in Appendix B), we obtain, for each $j = 2, \dots, \log^*(n/\delta)$:

$$|\mathcal{V}| \leq (d_0 + 1) \mathbf{Exp}_{I_j} \left[|\mathcal{V}_{|_{I_j}}| \right]. \quad (4)$$

We note that since the bounds (1)–(4) involve a dependency on the VC-dimension d_0 , we will sometimes need to explicitly refer to this parameter in addition to the primal shatter dimension d . Nevertheless, throughout the analysis we exploit the relation $d \leq d_0 = O(d \log d)$, mentioned in Section 1.

3 The Analysis: Refining Haussler’s Approach

Overview of the approach. We next present the proof of Theorem 1.4. In what follows, we assume that \mathcal{V} is δ -separated. We first recall the assumption that the primal shatter function of \mathcal{V} has a (d, d_1) Clarkson-Shor property, and that the length of each vector $\mathbf{v} \in \mathcal{V}$ under the L^1 norm is most k . This implies that \mathcal{V} consists of at most $O(n^{d_1} k^{d-d_1})$ vectors.

Since the Clarkson-Shor property is hereditary, then this also applies to any projection of \mathcal{V} onto a subset of indices, implying that the bound on $|\mathcal{V}_{|_I}|$ is at most $O(m^{d_1} k^{d-d_1})$, where I is a subset of $m - 1$ indices as above. However, due to our sampling scheme we expect that the length of each vector in $\mathcal{V}_{|_I}$ should be much smaller than k , (e.g., in expectation this value should not exceed $k(m-1)/n$), from which we may conclude that the actual bound on $|\mathcal{V}_{|_I}|$ is smaller than the trivial bound $O(m^{d_1} k^{d-d_1})$. Ideally, we would like to show that this bound is $O(m^{d_1} (km/n)^{d-d_1}) = O(n^{d_1} k^{d-d_1} / \delta^d)$, which matches our asymptotic bound in Theorem 1.4 (recall that $m = O(n/\delta)$). However, this is likely to happen only in case where the length of each vector in $\mathcal{V}_{|_I}$ does not exceed its expected value, or that there are only a few vectors whose length deviates from its expected value by far, whereas, in the worst case there might be many leftover “long” vectors in $\mathcal{V}_{|_I}$. Nevertheless, our goal is to show that, with some carefulness one can proceed in iterations, where initially I is a slightly larger sample, and then at each iteration we reduce its size, until eventually it becomes $O(m)$ and we remain with only a few long vectors. At each such iteration $\mathcal{V}_{|_I}$ is a random structure that depends on the choice of I and may thus contain long vectors, however, in expectation they will be scarce!

Specifically, we proceed over at most $\log^*(n/\delta)$ iterations, where we perform local improvements over the bound on $|\mathcal{V}|$, as follows. Let $|\mathcal{V}|^{(j)}$ be the bound on $|\mathcal{V}|$ after the j th iteration is completed, $1 \leq j \leq$

$\log^*(n/\delta)$. We first show in Corollary 3.2 that for the first iteration, $|\mathcal{V}| \leq |\mathcal{V}|^{(1)} = O\left(\frac{n^{d_1} k^{d-d_1} \log^d(n/\delta)}{\delta^d}\right)$, with a constant of proportionality that depends on d . Then, at each further iteration $j \geq 2$, we select a set I_j of $m_j - 1 = O(n \log^{(j)}(n/\delta)/\delta)$ indices uniformly at random without replacements from $[n]$ (see (3) for the bound on m_j). Our goal is to bound $\mathbf{Exp}_{I_j} [|\mathcal{V}_{|I_j}|]$ using the bound $|\mathcal{V}|^{(j-1)}$, obtained at the previous iteration, which, we assume by induction to be $O\left(\frac{n^{d_1} k^{d-d_1} (\log^{(j-1)}(n/\delta))^d}{\delta^d}\right)$ (where the base case $j = 2$ is shown in Corollary 3.2).

A key property in the analysis is then to show that the probability that the length of a vector $\mathbf{v} \in \mathcal{V}_{|I_j}$ (after the projection of \mathcal{V} onto I_j) deviates from its expectation decays exponentially (Lemma 3.3). Note that in our case this expectation is at most $k(m_j - 1)/n$. This, in particular, enables us to claim that *in expectation* the overall majority of the vectors in $\mathcal{V}_{|I_j}$ have length at most $O(k(m_j - 1)/n)$, whereas the remaining longer vectors are scarce. Specifically, since the Clarkson-Shor property is hereditary, we apply it to $\mathcal{V}_{|I_j}$ and conclude that the number of its vectors of length at most $O(k(m_j - 1)/n)$ is only $O\left(\frac{n^{d_1} k^{d-d_1} (\log^{(j)}(n/\delta))^d}{\delta^d}\right)$, with a constant of proportionality that depends on d . On the other hand, due to Lemma 3.3 and our inductive hypothesis, the number of longer vectors does not exceed $O\left(\frac{n^{d_1} k^{d-d_1}}{\delta^d}\right)$, which is dominated by the first bound. We thus conclude $\mathbf{Exp}_{I_j} [|\mathcal{V}_{|I_j}|] = O\left(\frac{n^{d_1} k^{d-d_1} (\log^{(j)}(n/\delta))^d}{\delta^d}\right)$. Then we apply Inequality (4) in order to complete the inductive step, whence we obtain the bound on $|\mathcal{V}|^{(j)}$, and thus on $|\mathcal{V}|$. These properties are described more rigorously in Lemma 3.4, where derive a recursive inequality for $|\mathcal{V}|^{(j)}$ using the bound on $\mathbf{Exp}_{I_j} [|\mathcal{V}_{|I_j}|]$. We emphasize the fact that the sample I_j is always chosen from the *original* ground set $[n]$, and thus, at each iteration we construct a new sample *from scratch*, and then exploit our observation in (4).

3.1 The First Iteration

In order to show our bound on $|\mathcal{V}|^{(1)}$, we form a subset $I_1 = (i_1, \dots, i_{m_1})$ of $m_1 = |I_1| = O\left(\frac{dn \log(n/\delta)}{\delta}\right)$ indices³ with the following two properties: (i) each vector in \mathcal{V} is mapped to a distinct vector in $\mathcal{V}_{|I_1}$, and (ii) the length of each vector in $\mathcal{V}_{|I_1}$ does not exceed $O(k \cdot m_1/n)$.

Lemma 3.1. *A sample I_1 as above satisfies properties (i)–(ii), with probability at least $1/2$.*

A set I_1 as above exists by the considerations in [13]. Nevertheless, in Appendix C we present the proof of Lemma 3.1 for the sake of completeness and clarity.

We next apply Lemma 3.1 in order to bound $|\mathcal{V}_{|I_1}|$. We first recall that the (d, d_1) Clarkson-Shor property of the primal shatter function of \mathcal{V} is hereditary. Incorporating the bound on m_1 and property (ii), we conclude that

$$|\mathcal{V}_{|I_1}| = O\left(m_1^{d_1} \left(\frac{km_1}{n}\right)^{d-d_1}\right) = O\left(\frac{n^{d_1} k^{d-d_1} \log^d(n/\delta)}{\delta^d}\right),$$

³In this particular step we use a different machinery than that of Haussler [19]; see the proof of Lemma 3.1 and our remark after Corollary 3.2. Therefore, $|I_1| = m_1$, rather than $m_1 - 1$. Furthermore, the constant of proportionality in the bound on m_1 depends just on the primal shatter dimension d instead of the VC-dimension d_0 as in (3).

with a constant of proportionality that depends on d . Now, due to property (i), $|\mathcal{V}| \leq |\mathcal{V}_{|_{I_1}}|$, we thus conclude:

Corollary 3.2. *After the first iteration we have: $|\mathcal{V}| \leq |\mathcal{V}|^{(1)} = O\left(\frac{n^{d_1} k^{d-d_1} \log^d(n/\delta)}{\delta^d}\right)$, with a constant of proportionality that depends on d .*

Remark: We note that the preliminary bound given in Corollary 3.2 is crucial for the analysis, as it constitutes the base for the iterative process described in Section 3.2. In fact, this step of the analysis alone bypasses our refinement to Haussler’s approach, and instead exploits the approach of Dudley [11].

3.2 The Subsequent Iterations: Applying the Inductive Step

Let us now fix an iteration $j \geq 2$. As noted above, we assume by induction on j that the bound $|\mathcal{V}|^{(j-1)}$ on $|\mathcal{V}|$ after the $(j-1)$ th iteration is $O\left(\frac{n^{d_1} k^{d-d_1} (\log^{(j-1)}(n/\delta))^d}{\delta^d}\right)$. Let I_j be a subset of $m_j - 1$ indices, chosen uniformly at random without replacements from $[n]$, with m_j given by (3). Let $\mathbf{v} \in \mathcal{V}$, and denote by $\mathbf{v}_{|_{I_j}}$ its projection onto I_j . The expected length $\mathbf{Exp}[\|\mathbf{v}_{|_{I_j}}\|]$ of $\mathbf{v}_{|_{I_j}}$ is at most $k(m_j - 1)/n = O(d_0 k \log^{(j)}(n/\delta)/\delta)$. We next show (see Appendix D for the proof):

Lemma 3.3 (Exponential Decay Lemma).

$$\mathbf{Prob} \left[\|\mathbf{v}_{|_{I_j}}\| \geq t \cdot \frac{k(m_j - 1)}{n} \right] < 2^{-tk(m_j-1)/n},$$

where $t \geq 2e$ is a real parameter and e is the base of the natural logarithm.

We now proceed as follows. Recall that we assume $k \geq \delta/2$, and by (3) we have $m_j = O\left(\frac{d_0 n \log^{(j)}(n/\delta)}{\delta}\right)$. Thus it follows from Lemma 3.3 that

$$\mathbf{Prob} \left[\|\mathbf{v}_{|_{I_j}}\| \geq C \cdot \frac{k(m_j - 1)}{n} \right] < \frac{1}{(\log^{(j-1)}(n/\delta))^D}, \quad (5)$$

where $C \geq 2e$ is a sufficiently large constant, and $D > d_0$ is another constant whose choice depends on C and d_0 , and can be made arbitrarily large. Since $d_0 \geq d$ we obviously have $D > d$. We next show:

Lemma 3.4. *Under the assumption that $k \geq \delta/2$, we have, at any iteration $j \geq 2$:*

$$|\mathcal{V}|^{(j)} \leq A(d_0 + 1) \cdot \frac{n^{d_1} k^{d-d_1} (\log^{(j)}(n/\delta))^d}{\delta^d} + (d_0 + 1) \cdot \frac{|\mathcal{V}|^{(j-1)}}{(\log^{(j-1)}(n/\delta))^D}, \quad (6)$$

where $|\mathcal{V}|^{(l)}$ is the bound on $|\mathcal{V}|$ after the l th iteration, and $A > 0$ is a constant that depends on d (and d_0) and the constant of proportionality determined by the Clarkson-Shor property of \mathcal{V} .

Proof. We in fact show:

$$\mathbf{Exp}_{I_j} [|\mathcal{V}_{|_{I_j}}|] \leq A \cdot \frac{n^{d_1} k^{d-d_1} (\log^{(j)}(n/\delta))^d}{\delta^d} + \frac{|\mathcal{V}|^{(j-1)}}{(\log^{(j-1)}(n/\delta))^D},$$

and then exploit the relation $|\mathcal{V}| \leq (d_0 + 1) \mathbf{Exp}_{I_j} \left[|\mathcal{V}_{|_{I_j}}| \right]$ (Inequality (4)), in order to prove (6).

In order to obtain the first term in the bound on $\mathbf{Exp}_{I_j} \left[|\mathcal{V}_{|_{I_j}}| \right]$, we consider all vectors of length at most $C \cdot \frac{k(m_j-1)}{n}$ (where $C \geq 2e$ is a sufficiently large constant as above) in the projection of \mathcal{V} onto a subset I_j of $m_j - 1$ indices (in this part of the analysis I_j can be arbitrary). Since the primal shatter function of \mathcal{V} has a (d, d_1) Clarkson-Shor property, which is hereditary, we obtain at most

$$O(m_j^{d_1} (k(m_j - 1)/n)^{d-d_1}) = O \left(\frac{n^{d_1} k^{d-d_1} (\log^{(j)}(n/\delta))^d}{\delta^d} \right)$$

vectors in $\mathcal{V}_{|_{I_j}}$ of length smaller than $C \cdot \frac{k(m_j-1)}{n} = O(\frac{k \log^{(j)}(n/\delta)}{\delta})$. It is easy to verify that the constant of proportionality A in the bound just obtained depends on d, d_0 , and the constant of proportionality determined by the Clarkson-Shor property of \mathcal{V} .

Next, in order to obtain the second term, we consider the vectors $\mathbf{v} \in \mathcal{V}$ that are mapped to vectors $\mathbf{v}_{|_{I_j}} \in \mathcal{V}_{|_{I_j}}$ with $\|\mathbf{v}_{|_{I_j}}\| > C \cdot \frac{k(m_j-1)}{n}$. By Inequality (5):

$$\mathbf{Exp} \left[\left| \left\{ \mathbf{v} \in \mathcal{V} \mid \|\mathbf{v}_{|_{I_j}}\| > C \cdot \frac{k(m_j-1)}{n} \right\} \right| \right] < \frac{|\mathcal{V}|}{(\log^{(j-1)}(n/\delta))^D},$$

and recall that $|\mathcal{V}|^{(j-1)}$ is the bound on $|\mathcal{V}|$ after the previous iteration $j - 1$. This completes the proof of the lemma. □

Remark: We note that the bound on $\mathbf{Exp}_{I_j} \left[|\mathcal{V}_{|_{I_j}}| \right]$ consists of the *worst-case* bound on the number of short vectors of length at most $C \cdot k(m_j - 1)/n$, obtained by the Clarkson-Shor property, plus the *expected* number of long vectors.

Wrapping up. We now complete the analysis and solve Inequality (6). Our initial assumption that $\delta \leq n/2^{(d_0+1)}$, and the fact that $D > d$ is sufficiently large, imply that the coefficient of the recursive term is smaller than 1, for any $2 \leq j \leq 1 + \log^*(n/\delta) - \log^*(d_0 + 1)$.⁴ Then, using induction on j , one can verify that the solution is

$$|\mathcal{V}|^{(j)} \leq 2A(d_0 + 1) \frac{n^{d_1} k^{d-d_1} (\log^{(j)}(n/\delta))^d}{\delta^d}, \quad (7)$$

for any $2 \leq j \leq 1 + \log^*(n/\delta) - \log^*(d_0 + 1)$.

We thus conclude $|\mathcal{V}|^{(j)} = O \left(\frac{n^{d_1} k^{d-d_1} (\log^{(j)}(n/\delta))^d}{\delta^d} \right)$. In particular, at the termination of the last iteration $j^* = 1 + \log^*(n/\delta) - \log^*(d_0 + 1)$, we obtain:

$$|\mathcal{V}| \leq |\mathcal{V}|^{(j^*)} = O \left(\frac{n^{d_1} k^{d-d_1}}{\delta^d} \right),$$

with a constant of proportionality that depends on d (and d_0). This at last completes the proof of Theorem 1.4.

⁴We observe that $2 \leq 1 + \log^*(n/\delta) - \log^*(d_0 + 1) \leq \log^*(n/\delta)$, due to our assumption that $\delta \leq n/2^{(d_0+1)}$, and the fact that $d_0 \geq 1$.

4 Applications

4.1 Realization to Geometric Set Systems

We now incorporate the random sampling technique of Clarkson and Shor [10] with Theorem 1.4 in order to conclude that small shallow packings exist in several useful scenarios. This includes the case where \mathcal{V} represents: (i) A collection of halfspaces defined over an n -point set in d -space. In this case, for any integer parameter $0 \leq k \leq n$, the number of halfspaces that contain at most k points is $O(n^{\lfloor d/2 \rfloor} k^{\lceil d/2 \rceil})$, and thus the primal shatter function has a $(d, \lfloor d/2 \rfloor)$ Clarkson-Shor property. (ii) A collection of balls defined over an n -point set in d -space. Here, the number of balls that contain at most k points is $O(n^{\lfloor (d+1)/2 \rfloor} k^{\lceil (d+1)/2 \rceil})$, and therefore the primal shatter function has a $(d+1, \lfloor (d+1)/2 \rfloor)$ Clarkson-Shor property. (iii) A collection of *parallel slabs* (that is, each of these regions is enclosed between two parallel hyperplanes and has an arbitrary width), defined over an n -point set in d -space. The number of slabs, which contains at most k points is $O(n^d k)$. (iv) A dual set system of points in d -space and a collection F of n $(d-1)$ -variate (not necessarily continuous or totally defined) functions of *constant description complexity*. Specifically, the graph of each function is a semi-algebraic set in \mathbb{R}^d defined by a constant number of polynomial equalities and inequalities of constant maximum degree (see [34, Chapter 7] for a detailed description of these properties, which we omit here).⁵ In this case, \mathcal{V} is represented by the cells (of all dimensions) in the *arrangement* of the graphs of the functions in F (see Section 1 for the definition) that lie below at most k function graphs. This portion of the arrangement is also referred to as the *at most k -level*, and its combinatorial complexity is $O(n^{d-1+\varepsilon} k^{1-\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on d and ε . Thus the primal shatter function has a $(d, d-1+\varepsilon)$ Clarkson-Shor property.

All bounds presented in (i)–(iv) are well known in the field of computational geometry; we refer the reader to [10, 26, 34] for further details. We thus conclude:

Corollary 4.1. *Let $\mathcal{V} \subseteq \{0, 1\}^n$ be a set of indicator vectors representing a set system of halfspaces defined over an n -point set in d -space, and let δ, k be two integer parameters as in Theorem 1.4. Then:*

$$\mathcal{M}(\delta, k, \mathcal{V}) = O\left(\frac{n^{\lfloor d/2 \rfloor} k^{\lceil d/2 \rceil}}{\delta^d}\right),$$

where the constant of proportionality depends on d .

Corollary 4.2. *Let $\mathcal{V} \subseteq \{0, 1\}^n$ be a set of indicator vectors representing a set system of balls defined over an n -point set in d -space, and let δ, k be two integer parameters as in Theorem 1.4. Then:*

$$\mathcal{M}(\delta, k, \mathcal{V}) = O\left(\frac{n^{\lfloor (d+1)/2 \rfloor} k^{\lceil (d+1)/2 \rceil}}{\delta^{d+1}}\right),$$

where the constant of proportionality depends on d .

Corollary 4.3. *Let $\mathcal{V} \subseteq \{0, 1\}^n$ be a set of indicator vectors representing a set system of parallel slabs defined over an n -point set in d -space, and let δ, k be two integer parameters as in Theorem 1.4. Then:*

$$\mathcal{M}(\delta, k, \mathcal{V}) = O\left(\frac{n^d k}{\delta^{d+1}}\right),$$

where the constant of proportionality depends on d .

⁵In [34] it is also required that the projection of each function onto the plane $x_d = 0$ has a constant description complexity.

Corollary 4.4. *Let $\mathcal{V} \subseteq \{0, 1\}^n$ be a set of indicator vectors representing a dual set system of n $(d - 1)$ -variate (not necessarily continuous or totally defined) functions of constant description complexity and points in d -space. Let δ, k be two integer parameters as in Theorem 1.4. Then:*

$$\mathcal{M}(\delta, k, \mathcal{V}) = O\left(\frac{n^{d-1+\varepsilon} k^{1-\varepsilon}}{\delta^d}\right),$$

where the constant of proportionality depends on d and on ε .

4.2 Spanning Trees of Low Total Conflict Number

Suppose we are given a set X of n points in d -space and a set Σ of m regions defined over X . With a slight abuse of notation, we also refer to Σ as the corresponding set system defined over X .⁶ We now assume that the set system (X, Σ) is shallow and has a (d, d_1) Clarkson-Shor property. The question at hand is to construct a spanning tree over Σ , whose overall *conflict number* is small. This notion is related to *spanning trees with low crossing number*, defined by Welzl [37] (see also [9]), where the tree spans X (rather than Σ), in which case it has a geometric realization, where the edges of the tree are the line segments connecting points of X . Our structure is dual to that of Welzl [37] and defined as follows.

Let $G = (\Sigma, E)$ be a graph with vertex set Σ . We say that a point $x \in X$ *conflicts* with an edge $\{S, S'\} \in E$ if $x \in S \triangle S'$. We then define the *conflict number* of an edge e of G as the number of points $x \in X$ with which it is in conflict, and then the total conflict number of G is the sum of the conflict numbers, over all its edges. Using similar arguments as in [29, Lemma 5.18] and [37], one can show (we omit the easy proof):

Lemma 4.5. *Let Σ be a set system of m sets, defined over an n -point set X . Assume Σ has a (d, d_1) Clarkson-Shor property, and that $|S| \leq k$, for each $S \in \Sigma$, where $1 \leq k \leq n$ is an integer parameter. Then there exists a spanning tree \mathcal{T} over Σ , whose total conflict number is $O(n^{d_1/d} k^{1-d_1/d} m^{1-1/d})$.*

Constructing an approximation of the tree. In order to construct \mathcal{T} efficiently, we relax this problem to only *approximating* the spanning tree, and then use the machinery of Har-Peled and Indyk [16] in order to conclude that when Σ is a set system as in Corollaries 4.1–4.3, a $(1 + \eta)$ -factor approximation for \mathcal{T} (that is, a spanning tree whose total conflict number is at most $(1 + \eta)$ of the smallest such number) can be computed in subquadratic time, for any $\eta > 0$. We present a sketch of this construction in Appendix E, and conclude:

Corollary 4.6. *Let X be a set of n points in d -space and let Σ be a set of m regions defined over X . Then, for any $\eta > 0$, one can compute in time*

- (i) $O^*\left(n^{\frac{d}{d+1}} m^{\frac{d}{d+1}} + n + m^{1+\frac{1}{1+\eta}}\right)$ a spanning tree of Σ , where Σ is a set of halfspaces in d -space, with overall $O\left((1 + \eta)n^{\frac{\lfloor d/2 \rfloor}{d}} k^{\frac{\lceil d/2 \rceil}{d}} m^{1-\frac{1}{d}}\right)$ conflicts.
- (ii) $O^*\left(n^{\frac{d+1}{d+2}} m^{\frac{d+1}{d+2}} + n + m^{1+\frac{1}{1+\eta}}\right)$, a spanning tree as above, where Σ is a set of balls in d -space, with overall $O\left((1 + \eta)n^{\frac{\lfloor (d+1)/2 \rfloor}{d+1}} k^{\frac{\lceil (d+1)/2 \rceil}{d+1}} m^{1-\frac{1}{d+1}}\right)$ conflicts.

⁶Here, Σ has the role \mathcal{V} in our original notation.

(iii) $O^* \left(n^{\frac{d}{d+1}} m^{\frac{d}{d+1}} + n + m^{1+\frac{1}{1+\eta}} \right)$, a spanning tree as above, where Σ is a set of parallel slabs in d -space, with overall $O \left((1+\eta)n^{1-1/(d+1)} k^{1/(d+1)} m^{1-\frac{1}{d+1}} \right)$ conflicts.

In the above bounds $O^*(\cdot)$ hides a poly-logarithmic factor.

Based on the above machinery, we next propose a general framework for updating the optimal solution (or an approximate solution) of a prescribed geometric optimization problem, over all regions in Σ .

A general framework. Let X, Σ be as in Lemma 4.5, and let $f : \Sigma \rightarrow \mathbb{R}$ be a function that assigns real values on the sets $S \in \Sigma$ (each being a subset of X). For example, f may correspond to diameter, width, radius of the smallest enclosing ball, volume of the minimum bounding box, etc., see, e.g., [3], where these measures are referred to as “faithful measures”. Our goal is to efficiently compute $f(S)$ for each $S \in \Sigma$.

Specifically, we assume to have a data-structure \mathcal{D} that maintains a subset $S \subseteq X$ with the following properties: (i) The time to preprocess \mathcal{D} is $P(n)$. (ii) The time to update \mathcal{D} (that is, inserting or deleting an element from \mathcal{D}) is $U(n)$. (iii) At any given time, querying \mathcal{D} for the value of f , w.r.t. the set K of the currently stored points, costs $Q(|K|)$ time.

Having this machinery at hand, in a brute-force approach, \mathcal{D} is initially empty. Then, for each $S \in \Sigma$, we insert its elements into \mathcal{D} , obtain $f(S)$ by querying \mathcal{D} , and then remove all these elements from \mathcal{D} . We proceed in this manner until all sets $S \in \Sigma$ are exhausted. Under the assumption that Σ is k -shallow, the resulting running time is $O(P(n) + mkU(n) + mQ(k))$. On the other hand, with the existence of a $(1+\eta)$ -factor approximation for the spanning tree \mathcal{T} (with properties as in Lemma 4.5), we can proceed as follows. With a slight abuse of notation, we also denote the approximate tree by \mathcal{T} . Initially, \mathcal{D} is empty as above, and we choose an arbitrary set $S \in \Sigma$, for which we compute $f(S)$ as above. Then we traverse \mathcal{T} from S in a BFS manner, update \mathcal{D} accordingly, and make a query at each vertex tracked during the search. Clearly, the number of these updates is proportional to the overall number of conflicts in \mathcal{T} , and thus the overall running time is $O \left(C(n, m, \eta) + P(n) + kU(n) + (1+\eta)n^{\frac{d}{d+1}} k^{1-\frac{d}{d+1}} m^{1-\frac{1}{d+1}} U(n) + mQ(k) \right)$, where $C(n, m, \eta)$ is the time to construct a $(1+\eta)$ -factor approximation for \mathcal{T} . We are interested in the scenario where this solution outperforms the brute-force algorithm (at least for some values of k). Below we describe a concrete scenario, related to the approximation of the faithful measures listed above, which demonstrates the usefulness of our framework.

Dynamic coresets. Based on the seminal work of Agarwal *et al.* [3] on *coresets*, Chan [7] presented a data structure, which maintains a constant-size coreset, with respect to “extent”, in $U(n) = O(\log n)$ update time, for all constant dimensions, with linear space and preprocessing time, where the constant of proportionality depends on the error parameter $\varepsilon > 0$. Using this machinery, it is straightforward to obtain dynamic $(1+\varepsilon)$ -factor approximation algorithms with logarithmic update time for computing the faithful measures stated above. The time $Q(k)$ to compute an approximation for these measures depends on ε and the dimension d . For simplicity of presentation, we omit the dependency on ε in the bounds of $U(n)$, $Q(k)$, and compare the performance of our approach w.r.t. the brute force computation, when we consider only the parameters n , m , and k .

With this machinery, the brute-force algorithm runs in $O(mk \log n + n)$ time, whereas our algorithm runs in $O \left(C(m, n, \eta) + k \log n + (1+\eta)n^{\frac{d}{d+1}} k^{1-\frac{d}{d+1}} m^{1-\frac{1}{d+1}} \log n + m + n \right)$ time. We now consider the bounds

stated in Corollary 4.6 and set $\eta = \log m$, in which case the term $O^*(m^{1+\frac{1}{1+\eta}})$ becomes nearly-linear and we pay only an extra logarithmic factor in the total number of conflicts. We then conclude:

Corollary 4.7. *Let X be a set of n points in d -space and a let Σ be a set of m regions defined over X , where, for each $S \in \Sigma$, $|S| \leq k$, where $1 \leq k \leq n$ is an integer parameter. Then one can compute a $(1 + \varepsilon)$ -factor approximation for the aforementioned faithful measures in time that is the minimum of $O(mk \log n + n)$ and*

- (i) $O^*\left(n^{\frac{d}{d+1}} m^{\frac{d}{d+1}} + k \log n + n^{\lfloor \frac{d/2 \rfloor} d} k^{\lfloor \frac{d/2 \rfloor} d} m^{1-\frac{1}{d}} + m + n\right)$, if Σ is a collection of halfspaces, in d -space.
- (ii) $O^*\left(n^{\frac{d+1}{d+2}} m^{\frac{d+1}{d+2}} + k \log n + n^{\lfloor \frac{(d+1)/2 \rfloor}{d+1}} k^{\lfloor \frac{(d+1)/2 \rfloor}{d+1}} m^{1-\frac{1}{d+1}} + m + n\right)$, if Σ is a collection of balls in d -space.
- (iii) $O^*\left(n^{\frac{d}{d+1}} m^{\frac{d}{d+1}} + k \log n + n^{\frac{d}{d+1}} k^{\frac{1}{d+1}} m^{1-\frac{1}{d+1}} + m + n\right)$, if Σ is a collection of parallel slabs in d -space.

The constant of proportionality in each of these bounds depends on ε and d .

4.3 Geometric Discrepancy

Given a set system (X, Σ) as above, we now wish to color the points of X by two colors, such that in each set of Σ the deviation from an even split is as small as possible.

Formally, a *two-coloring* of X is a mapping $\chi : X \rightarrow \{-1, +1\}$. For a set $S \in \Sigma$ we define $\chi(S) := \sum_{x \in S} \chi(x)$. The *discrepancy* of Σ is then defined as $\text{disc}(\Sigma) := \min_{\chi} \max_{S \in \Sigma} |\chi(S)|$.

In a previous work [13], the author presented size-sensitive discrepancy bounds for set systems of halfspaces defined over n points in d -space. These bounds were achieved by combining the *entropy method* [24] with δ -packings, and, as observed in [13], they are optimal up to a poly-logarithmic factor. Incorporating our bound in Theorem 1.4 into the analysis in [13], the bounds on $\chi(S)$ improve by a $\sqrt{\log n}$ factor. Specifically, we obtain (we omit the technical details in this version):

Corollary 4.8. *Let Σ be a set system of halfspaces defined over n -points in d -space ($d \geq 3$). Then, there is a two-coloring χ , such that for each $S \in \Sigma$, $\chi(S) = O\left(|S|^{1/4} n^{1/4-1/(2d)} \log^{1/2d} n\right)$, for $d \geq 4$ even, and $\chi(S) = O\left(|S|^{1/4+1/(4d)} n^{1/4-3/(4d)} \log^{1/2d} n\right)$, for $d \geq 5$ odd, where the constant of proportionality depends on d . When $d = 3$ the bound on $\chi(S)$ is $O\left(|S|^{1/3} \log^{7/6} n\right)$.⁷*

5 Concluding Remarks and Further Research

We note that Corollary 4.4 implies that one can pack the “shallow level” in an arrangement of $(d - 1)$ -variate function graphs F (as defined in Section 4.1) with a relatively small number of Hamming balls. In fact, if those functions are just hyperplanes, then by a standard point-hyperplane duality, Corollary 4.1 implies that the number of Hamming balls that pack the at most k -level in the underlying arrangement is

⁷We note that the case $d = 2$ in Corollary 4.8 has been resolved by Har-Peled and Sharir [17].

$O(n^{\lfloor d/2 \rfloor} k^{\lceil d/2 \rceil} / \delta^d)$. Roughly speaking, this particular bound can be obtained by either (i) *shallow cuttings* [27] and a standard reduction between packing and covering. Roughly speaking, this is a coverage of the shallow level of the arrangement by a collection Ξ of (possibly unbounded) pairwise-disjoint “primitive” cells, such that the interior of each cell is crossed by a small fraction of the function graphs in F . When F is a collection of hyperplanes the packing number and the bound on $|\Xi|$ are asymptotically the same (where each cell of Ξ corresponds to a “center” of a Hamming ball). Or (ii) the bound $\Omega(\delta^d)$ on the “volume” of a Hamming ball of radius δ (as observed in [37]) and the fact that the complexity of the at most k -level in an arrangement of n hyperplanes is $O(n^{\lfloor d/2 \rfloor} k^{\lceil d/2 \rceil})$ [10].⁸ Nevertheless, these approaches may not be applicable when the input functions are more general (e.g., satisfy the properties stated in Section 4.1 for, say, $d \geq 4$), in this case one may use *linearization* but that may result in an overestimated bound. We hope our bound on the packing number will be useful in geometric computing, and, in particular in the context of spanning trees of low (total) crossing number.

The analogy between shallow packings for dual set systems and shallow cuttings may cause one to interpret shallow packings as the “primal version” of shallow cuttings. In this paper we named two useful applications for shallow packings. We hope to find additional applications in geometry and beyond.

Acknowledgments. The authors wishes to thank Boris Aronov, Sarel Har-Peled, Aryeh Kontorovich, and Wolfgang Mulzer for useful discussions and suggestions. In particular, the author is grateful to Sarel Har-Peled for suggesting the application described in Section 4.2. Last but not least, the author deeply thanks Ramon Van Handel, for various discussions and for spotting an error in an earlier version of this paper. In particular, work on this paper began due to a discussion with Ramon after a talk the author gave in Princeton University.

References

- [1] P. K. Agarwal, A. Efrat, and M. Sharir. Vertical decomposition of shallow levels in 3-dimensional arrangements and its applications. *SIAM J. Comput.*, 29(2000):912–953.
- [2] P. K. Agarwal and J. Erickson. Geometric range searching and its relatives. *Discrete Comput. Geom.* (1997).
- [3] P. K. Agarwal, S. Har-Peled, and K. R. Varadarajan. Approximating extent measures of Ppoints. *J. ACM*, 51(4):606–635 (2004).
- [4] N. Alon and J. H. Spencer. *The Probabilistic Method*. 2nd Edition, Wiley-Interscience, New York, USA, 2000.
- [5] A. Auger and B. Doerr. *Theory of Randomized Search Heuristics: Foundations and Recent Developments*, World Scientific Publishing, 2011.
- [6] N. H. Bshouty, Y. Li, and P. M. Long. Using the doubling dimension to analyze the generalization of learning algorithms. *J. Comput. System Sci.*, 75(6):323–335 (2009).

⁸We omit the official definition of shallow cuttings and the straightforward details in this discussion, and refer the non-expert reader to [1, 27, 26, 37] and the references therein.

- [7] T. M. Chan. Dynamic coresets. *Discrete Comput. Geom.*, 42: 469–488 (2009).
- [8] B. Chazelle. A note on Haussler's packing lemma. Unpublished manuscript, Princeton (1992).
- [9] B. Chazelle and E. Welzl. Quasi-optimal range searching in spaces of finite VC-dimension. *Discrete Comput. Geom.*, 4:467–489 (1989).
- [10] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421 (1989).
- [11] R. M. Dudley. Central limit theorems for empirical measures. *Ann. Probab.*, 6(6):899–1049 (1978).
- [12] D. Eppstein. Dynamic Euclidean minimum spanning trees and extrema of binary functions. *Discrete Comput. Geom.*, 13:111–122 (1995).
- [13] E. Ezra. A size-sensitive discrepancy bound for set systems of bounded primal shatter dimension. In *Proc. Twenty-Fifth Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 1378–1388 (2014).
- [14] L. Gottlieb, A. Kontorovich, and E. Mossel. VC bounds on the cardinality of nearly orthogonal function classes. *Discrete Math.*, 312(10):1766–1775 (2012).
- [15] S. Har-Peled. *Geometric Approximation Algorithms*, Mathematical Surveys and Monographs, Vol. 173 (2011).
- [16] S. Har-peled and P. Indyk. When crossings count - approximating the minimum spanning tree. In *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, pp. 166–175 (2002).
- [17] S. Har-Peled and M. Sharir, Relative (p, ε) -approximations in geometry, *Discrete Comput. Geom.*, 45(3):462–496 (2011).
- [18] D. Haussler. Decision theoretic generalizations of the PAC model for neural net and other learning applications. In *Information and Computation*, 100(1):78–150 (1992).
- [19] D. Haussler. Sphere packing numbers for subsets of the Boolean n -cube with bounded Vapnik-Chervonenkis dimension. *J. Combinatorial Theory Ser. A*, 69:217–232 (1995).
- [20] D. Haussler, N. Littlestone, M. K. Warmuth. Predicting $\{0, 1\}$ -functions on randomly drawn points. *Information and Computation*, 115(2), 248–292 (1994).
- [21] D. Haussler and E. Welzl. ε -nets and simplex range queries. *Discrete Comput. Geom.*, 2:127–151 (1987).
- [22] P. Indyk and R. Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proc. 30th Annu. ACM Sympos. Theory Comput.*, pp. 604–613 (1998).
- [23] Y. Li, P. M. Long, and A. Srinivasan. Improved bounds on the sample complexity of learning. *J. Comput. Sys. Sci.*, 62(3):516–527 (2001).
- [24] S. Lovett and R. Meka. Constructive discrepancy minimization by walking on the edges. In *Proc. 53th Annu. IEEE Symp. Found. Comput. Sci.*, 61–67, (2012).

- [25] E. Kushilevitz, R. Ostrovsky, and Y. Rabani. Efficient search for approximate nearest neighbor in high dimensional spaces. *SIAM J. Comput.*, 2(30):457–474 (2000).
- [26] J. Matoušek, *Lectures on Discrete Geometry*, Springer-Verlag New York (2002).
- [27] J. Matoušek. Reporting points in halfspaces. *Comput. Geom. Theory Appl.*, 2:169–186 (1992).
- [28] J. Matoušek. Tight upper bounds for the discrepancy of halfspaces. *Discrete Comput. Geom.*, 13:593–601 (1995).
- [29] J. Matoušek. *Geometric Discrepancy, Algorithms and Combinatorics*, Vol. 18, Springer Verlag, Heidelberg (1999).
- [30] W. Mulzer. Chernoff Bounds, *Personal note*.
<http://page.mi.fu-berlin.de/mulzer/notes/misc/chernoff.pdf>.
- [31] A. Panconesi and A. Srinivasan. Randomized distributed edge coloring via an extension of the Chernoff-Hoeffding bounds. *SIAM J. Comput.*, 26:350–368 (1997).
- [32] D. Pollard. *Convergence of Stochastic Processes*, Springer-Verlag (1984).
- [33] N. Sauer. On the density of families of sets. *J. Combin. Theory, Ser A*, 13(1): 145–147 (1972).
- [34] M. Sharir and P. K. Agarwal. *Davenport-Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, New York (1995).
- [35] S. Shelah. A combinatorial problem, stability and order for models and theories in infinitary languages. *Pacific J. Math.*, 41:247–261 (1972).
- [36] V. Vapnik and A. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory Prob. Appl.*, 16(2):264–280 (1971).
- [37] E. Welzl. On spanning trees with low crossing numbers. In *Data Structures and Efficient Algorithms, Final Report on the DFG Special Joint Initiative*, volume 594 of *Lect. Notes in Comp. Sci.*, Springer-Verlag, Heidelberg, pp. 233–249 (1992).

A An Example of a Shallow Set System with Large Packing Numbers

Continuing our construction from Section 1, the ground set is a collection of axis-parallel rectangles, and the vectors \mathcal{V} represent subsets of rectangles that cover a point in the plane. For simplicity of exposition, we define these vectors to represent all two-dimensional cells in the arrangement of the given rectangles. It is well known that the primal shatter function of \mathcal{V} is quadratic (see, e.g., [15]), and therefore, by Theorem 1.3, the packing number is $O((n/\delta)^2)$. Nevertheless, we claim that even when the arrangement is shallow, the asymptotic bound on the packing number is not any better than $(n/\delta)^2$. Indeed, fix a positive even parameter $\delta > 0$, and suppose, without loss of generality, that n/δ is an integer number. Consider now an $\frac{n}{\delta} \times \frac{n}{\delta}$ grid of long and skinny rectangles, where each rectangle in the grid is duplicated $\delta/2$ times (with a possibly infinitesimal perturbation), as illustrated in Figure 1. Clearly, each (two-dimensional) cell in the arrangement is covered by at most δ rectangles, and thus the set system is δ -shallow. Consider now only




Figure 1: A grid of $\frac{n}{\delta} \times \frac{n}{\delta}$ axis-parallel rectangles, each of which with multiplicity $\delta/2$. The multiplicity in the figure is depicted only for the leftmost vertical and the top horizontal rectangles. The small shaded rectangle in the figure is a δ -shallow cell of the arrangement.

the cells at “depth” δ (that is, they are covered by precisely δ rectangles), and let $\mathcal{F} \subset \mathcal{V}$ be the set of their representative vectors. It is easy to verify that, for each pair $v, v' \in \mathcal{F}$, the distance $\rho(v, v')$ is at least δ (see once again Figure 1), and thus \mathcal{F} is both δ -separated and δ -shallow. However, by construction, we have $|\mathcal{F}| = \Omega((n/\delta)^2)$, and thus the δ -shallowness assumption does not yield an improvement over the general case.

B Overview of Haussler’s Approach

Let $\mathcal{V} \subseteq \{0, 1\}^n$ be a collection of indicator vectors of primal shatter dimension d . We denote its VC-dimension by d_0 ; as discussed in Section 1 $d_0 = O(d \log d)$.

We first form a probability distribution P over \mathcal{V} , implying that \mathcal{V} can be viewed⁹ as an n -dimensional random variable taking values in $\{0, 1\}^n$. Thus its components \mathcal{V}_i , $i = 1, \dots, n$, represent n correlated indicator random variables (Bernulli random variables), and each of their values is determined by randomly selecting a vector $\mathbf{v} \in \mathcal{V}$, and letting \mathcal{V}_i be the i th component of \mathbf{v} . The variance of a Bernulli random variable B is known to be $\mathbf{Prob}[B = 1] \mathbf{Prob}[B = 0]$, and then, for a sequence B_1, \dots, B_m of Bernulli random variables, the *conditional variance* of B_m given B_1, \dots, B_{m-1} is defined as

$$\mathbf{Var}(B_m | B_1, \dots, B_{m-1}) = \sum_{\mathbf{v} \in \{0, 1\}^{m-1}} \mathbf{Prob}(\mathbf{v}) \mathbf{Prob}(B_m = 1 | \mathbf{v}) (1 - \mathbf{Prob}(B_m = 1 | \mathbf{v})),$$

where $\mathbf{Prob}(\mathbf{v}) = \mathbf{Prob}(B_1 = \mathbf{v}_1, B_2 = \mathbf{v}_2, \dots, B_{m-1} = \mathbf{v}_{m-1})$, and $\mathbf{Prob}(B_m = 1 | \mathbf{v}) = \mathbf{Prob}(B_m = 1 | B_1 = \mathbf{v}_1, B_2 = \mathbf{v}_2, \dots, B_{m-1} = \mathbf{v}_{m-1})$.

A key property in the analysis of Haussler [19] lies in the density of a *unit distance graph* $G = (\mathcal{V}, E)$, defined over \mathcal{V} , whose edges correspond to all pairs $\mathbf{u}, \mathbf{v} \in \mathcal{V}$, whose symmetric difference distance is (precisely) 1. In other words, \mathbf{u}, \mathbf{v} appear as neighbors on the unit cube $\{0, 1\}^n$. It has been shown by Haussler *et al.* [20] that the density of G is bounded by the VC-dimension of \mathcal{V} , that is, $|E|/|\mathcal{V}| \leq d_0$;

⁹For the time being, P is an arbitrary distribution, but later on (Lemma B.2) it is taken to be the uniform distribution in the obvious way, where each vector in \mathcal{V} is equally likely to be chosen. This distribution, however, may not remain uniform after the projection of \mathcal{V} onto a proper subsequence $I' = (i_1, \dots, i_m)$ of $m < n$ indices, as several vectors in \mathcal{V} may be projected onto the same vector in $\mathcal{V}_{I'}$.

see also [19] for an alternative proof using the technique of “shifting”.¹⁰ Then, this low density property is exploited in order to show that once we have chosen $(n - 1)$ coordinates of the random variable \mathcal{V} , the variance in the choice of the remaining coordinate is relatively small. That is:

Lemma B.1 ([19]). *For any distribution P on \mathcal{V} ,*

$$\sum_{i=1}^n \mathbf{Var}(\mathcal{V}_i | \mathcal{V}_1, \dots, \mathcal{V}_{i-1}, \mathcal{V}_{i+1}, \dots, \mathcal{V}_n) \leq d_0.$$

As observed in [19], Lemma B.1 continues to hold on any restriction of \mathcal{V} to a sequence $I' = \{i_1, \dots, i_m\}$ of $m \leq n$ indices. Indeed, when projecting \mathcal{V} onto I' the VC-dimension in the resulting set system remains d_0 . Furthermore, the conditional variance is now defined w.r.t. the induced probability distribution on $\mathcal{V}_{|I'}$ in the obvious way, where the probability to obtain a sequence of m values corresponds to an appropriate marginal distribution, that is, $\mathbf{Prob}_{|I'}(u_1, \dots, u_m) = \mathbf{Prob}(\mathbf{v} \in \mathcal{V} \mid v_{i_j} = u_j, 1 \leq j \leq m)$. With this observation, we can rewrite the inequality stated in Lemma B.1 as

$$\sum_{i=1}^m \mathbf{Var}(\mathcal{V}_{i_j} | \mathcal{V}_{i_1}, \dots, \mathcal{V}_{i_{j-1}}, \mathcal{V}_{i_{j+1}}, \dots, \mathcal{V}_{i_m}) \leq d_0.$$

If I' is a sequence chosen uniformly at random (over all such m -tuples), then when averaging over all choices of I' we clearly obtain:

$$\mathbf{Exp} \left[\sum_{i=1}^m \mathbf{Var}(\mathcal{V}_{i_j} | \mathcal{V}_{i_1}, \dots, \mathcal{V}_{i_{j-1}}, \mathcal{V}_{i_{j+1}}, \dots, \mathcal{V}_{i_m}) \right] \leq d_0,$$

or

$$\sum_{i=1}^m \mathbf{Exp} [\mathbf{Var}(\mathcal{V}_{i_j} | \mathcal{V}_{i_1}, \dots, \mathcal{V}_{i_{j-1}}, \mathcal{V}_{i_{j+1}}, \dots, \mathcal{V}_{i_m})] \leq d_0,$$

by linearity of expectation. In fact, by symmetry of the random variables \mathcal{V}_{i_j} (recall that I' is a random m -tuple) each of the summands in the above inequality has an equal contribution, and thus, in particular (recall once again that the expectation is taken over all choices of $I' = \{i_1, \dots, i_m\}$):

$$\mathbf{Exp}_{I'} [\mathbf{Var}(\mathcal{V}_{i_m} | \mathcal{V}_{i_1}, \dots, \mathcal{V}_{i_{m-1}})] \leq \frac{d_0}{m}, \quad (8)$$

where we write $\mathbf{Exp}_{I'}[\cdot]$ to emphasize the fact that the expectation is taken over all choices of I' . The above bound is now integrated with the next key property:

Lemma B.2 ([19]). *Let \mathcal{V} be δ -separated subset of $\{0, 1\}^n$, for some $1 \leq \delta \leq n$ integer, and form a uniform distribution P on \mathcal{V} . Let $I = (i_1, \dots, i_{m-1})$ be a sequence of $m - 1$ distinct indices between 1 and n , where m is any integer between 1 and n . Suppose now that another index i_m is drawn uniformly at random from the remaining $n - m + 1$ indices. Then*

$$\mathbf{Exp} [\mathbf{Var}(\mathcal{V}_{i_m} | \mathcal{V}_{i_1}, \dots, \mathcal{V}_{i_{m-1}})] \geq \frac{\delta}{2(n - m + 1)} \left(1 - \frac{|\mathcal{V}_{|I}|}{|\mathcal{V}|} \right),$$

where the conditional variance is taken w.r.t. the distribution P , and the expectation is taken w.r.t. the random choice of i_m .

¹⁰We cannot guarantee such a relation when the VC-dimension d_0 is replaced by the primal shatter dimension d , and therefore we proceed with the analysis using this ratio.

We now observe that when the entire sequence $I' = (i_1, \dots, i_m)$ is chosen uniformly at random, then the bound in Lemma B.2 continues to hold when averaging on the entire sequence I' (rather than just on i_m), that is, we have:

$$\begin{aligned} \mathbf{Exp}_{I'} [\mathbf{Var} (\mathcal{V}_{i_m} | \mathcal{V}_{i_1} \dots, \mathcal{V}_{i_{m-1}})] &\geq \mathbf{Exp}_{I'} \left[\frac{\delta}{2(n-m+1)} \left(1 - \frac{|\mathcal{V}_{|I}|}{|\mathcal{V}|} \right) \right] \\ &= \frac{\delta}{2(n-m+1)} \left(1 - \frac{\mathbf{Exp}_I [|\mathcal{V}_{|I}|]}{|\mathcal{V}|} \right). \end{aligned} \quad (9)$$

Note that under this formulation, $|\mathcal{V}_{|I}|$ (the number of sets in the projection of \mathcal{V} onto I) is a random variable that depends on the choice of $I = (i_1, \dots, i_{m-1})$, in particular, since it does not depend on the choice of i_m , we have $\mathbf{Exp}_{I'} [|\mathcal{V}_{|I}|] = \mathbf{Exp}_I [|\mathcal{V}_{|I}|]$.

The analysis of Haussler [19] then proceeds as follows. We assume that \mathcal{V} is δ -separated as in Lemma B.2 (then a bound on $|\mathcal{V}|$ is the actual bound on the packing number), and then choose

$$m := \left\lceil \frac{(2d_0 + 2)(n + 1)}{\delta + 2d_0 + 2} \right\rceil,$$

indices i_1, \dots, i_m uniformly at random without replacements from $[n]$ (without loss of generality, we can assume that $\delta \geq 3$ as, otherwise, we set the bound on the packing to be $O(n^{d_1} k^{d-d_1})$, as asserted by the (d, d_1) Clarkson-Shor property of \mathcal{V} . Moreover, we can assume, without loss of generality, $n \geq d_0, \delta$, and thus we have $m \leq n$). Put $I' = \{i_1, \dots, i_m\}$, $I = I' \setminus \{i_m\}$. Then the analysis in [19] combines the two bounds in Inequalities (8) and (9) in order to derive an upper bound on $|\mathcal{V}|$, from which the bound in the Packing Lemma (Theorem 1.3) is obtained. Specifically, using simple algebraic manipulations, we obtain:

$$|\mathcal{V}| \leq \frac{\mathbf{Exp}_I [|\mathcal{V}_{|I}|]}{1 - \frac{2d_0(n-m+1)}{m\delta}} \quad (10)$$

It has been shown in [19] that due to our choice of m , we have:

$$\frac{2d_0(n-m+1)}{m\delta} \leq \frac{d_0}{d_0 + 1}, \quad (11)$$

from which we obtain Inequality (1), as asserted. This completes the description of our first extension of Haussler's analysis, as described in Section 2.

In order to apply our second extension, we observe that one only needs to assume $m \leq n$ when obtaining Inequalities (8) and (9). In addition, the choice of m in Inequality (10) can be made slightly larger, since the term $\frac{2d_0(n-m+1)}{m\delta}$ in Inequality (11) is a decreasing function of m , as can easily be verified. Recall that in our analysis we replace m by $m_j := m \log^{(j)}(n/\delta)$, where $2 \leq j \leq \log^*(n/\delta)$, in which case we still obtain $|\mathcal{V}| \leq (d_0 + 1) \mathbf{Exp}_{I_j} [|\mathcal{V}_{|I_j}|]$, where I_j is a set of $m_j - 1$ indices chosen uniformly at random without replacements from $[n]$.

C The First Iteration

Our proof relies on the notions of “relative approximations” and “ ε -nets”, define below:

Relative approximations and ε -nets. We mentioned in the introduction the notion of *relative (ε, η) -approximations*. We now define them formally: Following the definition from [17], given a set system $\mathcal{V} \subseteq \{0, 1\}^n$ and two parameters, $0 < \varepsilon < 1$ and $0 < \eta < 1$, we say that a subsequence I of indices is a *relative (ε, η) -approximation* if it satisfies, for each vector $\mathbf{v} \in \mathcal{V}$,

$$\left| \frac{\|\mathbf{v}_{|I}\|}{|I|} - \frac{\|\mathbf{v}\|}{n} \right| \leq \eta \frac{\|\mathbf{v}\|}{n}, \quad \text{if } \frac{\|\mathbf{v}\|}{n} \geq \varepsilon, \quad \text{and}$$

$$\left| \frac{\|\mathbf{v}_{|I}\|}{|I|} - \frac{\|\mathbf{v}\|}{n} \right| \leq \eta \varepsilon, \quad \text{otherwise,}$$

where $\mathbf{v}_{|I}$ is the projection of $\mathbf{v} \in \mathcal{V}$ onto I .

As observed by Har-Peled and Sharir [17], the analysis of Li *et al.* [23] implies that if \mathcal{V} has primal shatter dimension d , then a random sample of $\frac{cd \log(1/\varepsilon)}{\varepsilon \eta^2}$ indices (each of which is drawn independently) is a relative (ε, η) -approximation for \mathcal{V} with constant probability, where $c > 0$ is an absolute constant. More specifically, success with probability at least $1 - q$ is guaranteed if one samples $\frac{c(d \log(1/\varepsilon) + \log(1/q))}{\varepsilon \eta^2}$ indices.¹¹

It was also observed in [17] that ε -nets arise as a special case of relative (ε, η) -approximations. Specifically, an ε -net is a subsequence of indices I with the property that any vector $\mathbf{v} \in \mathcal{V}$ with $\|\mathbf{v}\| \geq n\varepsilon$ satisfies $\|\mathbf{v}_{|I}\| \geq 1$. In other words, I is a hitting set for all the “long” vectors. In this case, if we set η to be some constant fraction, say, $1/4$, then a relative $(\varepsilon, 1/4)$ -approximation becomes an ε -net. Moreover, a random sample of $O\left(\frac{d \log(1/\varepsilon) + \log(1/q)}{\varepsilon}\right)$ indices (with an appropriate choice of the constant of proportionality) is an ε -net for \mathcal{V} , with probability at least $1 - q$; see [17] for further details.

Proof of Lemma 3.1. In order to show (i), we first form the set system corresponding to all symmetric difference pairs induced by \mathcal{V} . That is, we form the vector set \mathcal{D} , where $\mathcal{D} = \{(|\mathbf{u}_1 - \mathbf{v}_1|, \dots, |\mathbf{u}_n - \mathbf{v}_n|) \mid \mathbf{u}, \mathbf{v} \in \mathcal{V}\}$. Since we assume that \mathcal{V} is δ -separated, we have $\|\mathbf{w}\| \geq \delta$, for each $\mathbf{w} \in \mathcal{D}$.

We now construct an ε -net for \mathcal{D} , with $\varepsilon = \delta/n$. By our discussion above a sample I_1 of $O(d(n/\delta) \log(n/\delta))$ indices has this property with probability greater than, say, $3/4$ (for a sufficiently large constant of proportionality). Thus, by definition, any vector $\mathbf{w} \in \mathcal{D}$ (recall that its length is at least δ) must satisfy $|\mathbf{w}_{|I_1}| \geq 1$, where $\mathbf{w}_{|I_1}$ denotes the projection of \mathbf{w} onto I_1 . But this implies that we must have $\mathbf{u}_{|I_1} \neq \mathbf{v}_{|I_1}$, for each pair $\mathbf{u}, \mathbf{v} \in \mathcal{V}$, and thus \mathbf{u}, \mathbf{v} must be mapped to distinct vectors in the projection of \mathcal{V} onto I_1 , from which property (i) follows.

In order to have property (ii) we observe that the same sample I_1 is also a *relative $(\delta/n, 1/4)$ -approximation* for \mathcal{V} with probability at least $3/4$ (for an appropriate choice of the constant of proportionality). Given this property of I_1 , this implies that any vector $\mathbf{v} \in \mathcal{V}$ satisfies

$$\left| \frac{\|\mathbf{v}_{|I_1}\|}{m_1} - \frac{\|\mathbf{v}\|}{n} \right| \leq \frac{1}{4} \cdot \frac{\|\mathbf{v}\|}{n},$$

¹¹We note that although in the original analysis for this bound d is the VC-dimension, this assumption can be replaced by having just a primal shatter dimension d ; see, e.g., [15] for the details of the analysis.

if $\|\mathbf{v}\|/n \geq \delta/n$, and

$$\left| \frac{\|\mathbf{v}_{|I_1}\|}{m_1} - \frac{\|\mathbf{v}\|}{n} \right| \leq \frac{1}{4} \cdot \frac{\delta}{n},$$

otherwise. Since $\|\mathbf{v}\| \leq k$, and $k \geq \delta/2$ by assumption, it is easy to verify that we always have $\|\mathbf{v}_{|I_1}\| \leq 3/2 \cdot km_1/n$. In other words, $\|\mathbf{v}_{|I_1}\| = O(k \cdot m_1/n)$, as asserted.

Combining the two roles of I_1 (each with probability $3/4$), it follows that it is both a (δ/n) -net for \mathcal{D} and a relative $(\delta/n, 1/4)$ -approximation for \mathcal{V} , with probability at least $1/2$, and thus it satisfies properties (i)–(ii) with this probability. This completes the proof of the lemma. \square

D Proof of Lemma 3.3

Proof. We first observe that the length of $\mathbf{v}_{|I_j}$ is a random variable with a hypergeometric distribution. Indeed, this is precisely the question of uniformly choosing $m_j - 1$ elements at random (into our set I_j) from a given set of n elements without replacements, and then, for a given $\|\mathbf{v}\|$ -element subset of the full set (recall that the length of \mathbf{v} corresponds to the cardinality of an appropriate subset in the set system), we consider how many of its elements have been chosen into I_j . Specifically, we have:

$$\mathbf{Prob} \left[\|\mathbf{v}_{|I_j}\| = s \right] = \frac{\binom{\|\mathbf{v}\|}{s} \binom{n-\|\mathbf{v}\|}{m_j-1-s}}{\binom{n}{m_j-1}},$$

for each non-negative integer $s \leq \min\{\|\mathbf{v}\|, m_j - 1\}$.

Our goal is to show a Chernoff-type bound over the probability that $\|\mathbf{v}_{|I_j}\|$ deviates from its expectation. However, we face the difficulty that the corresponding indicator variables are not independent, and thus we cannot apply a Chernoff bound directly (see, e.g. [4]). Nevertheless, in our scenario a Chernoff bound is still applicable, this can be viewed by various approaches, see, e.g., [5, 30, 31]. For the sake of completeness we describe the proof in detail, and rely on the analysis of Panconesi and Srinivasan [31], which implies that when the underlying indicator variables are “negatively correlated”, one can still apply a Chernoff bound (see also [5]).

We enumerate all non-zero coordinates of \mathbf{v} in an arbitrary order, let $L = \{l_1, \dots, l_{\|\mathbf{v}\|}\}$ be this set of indices (in this notation we ignore all the zero-coordinates), and attach an indicator variable X_i to each index $l_i \in L$, which is defined to be one if and only if $l_i \in I_j$ (in other words, the corresponding element in the underlying set induced by \mathbf{v} has been chosen to be included into the sample of the $m_j - 1$ elements). According to this notation, $\|\mathbf{v}_{|I_j}\|$ is represented by the sum $X = \sum_{i=1}^{\|\mathbf{v}\|} X_i$. It is now easy to verify that $\mathbf{Prob}[X_i = 1] = (m_j - 1)/n$, and by linearity of expectation $\mathbf{Exp}[\|\mathbf{v}_{|I_j}\|] = \|\mathbf{v}\| \cdot (m_j - 1)/n$.

However, the variables X_i are not independent due to our probabilistic model (that is, I_j is chosen without replacements), albeit, they are *negatively correlated*. This implies that for each subset $K \subseteq \{1, \dots, \|\mathbf{v}\|\}$

$$\mathbf{Prob} \left[\bigwedge_{i \in K} X_i = 0 \right] \leq \prod_{i \in K} \mathbf{Prob}[X_i = 0],$$

and

$$\mathbf{Prob} \left[\bigwedge_{i \in K} X_i = 1 \right] \leq \prod_{i \in K} \mathbf{Prob}[X_i = 1].$$

Indeed, following the considerations in [5], let us show first the latter inequality. Put $L_K = \bigcup_{i \in K} \{l_i\}$. Then $\mathbf{Prob}[\bigwedge_{i \in K} X_i = 1] = \mathbf{Prob}[L_K \subseteq I_j]$, and since I_j is uniformly chosen, in order to bound the latter we need to take the proportion between the number of subsets of size $m_j - 1$ that contain L_K and the entire number of subsets of size $m_j - 1$ that can be chosen from an n -element set. Hence

$$\mathbf{Prob}[L_K \subseteq I_j] = \frac{\binom{n-|K|}{m_j-1-|K|}}{\binom{n}{m_j-1}} = \frac{(m_j-1)(m_j-2) \cdots (m_j-|K|)}{n(n-1) \cdots (n-|K|+1)},$$

and the latter is smaller than $\left(\frac{m_j-1}{n}\right)^{|K|}$, as is easily verified. Using similar arguments for the first correlation inequality, we obtain

$$\mathbf{Prob}\left[\bigwedge_{i \in K} X_i = 0\right] = \frac{\binom{n-|K|}{m_j-1}}{\binom{n}{m_j-1}} < \left(1 - \frac{m_j-1}{n}\right)^{|K|}.$$

We are now ready to apply [31, Theorem 3.4] stating that if the indicator variables X_i are negatively correlated then (recall that $X = \sum_{i=1}^{\|\mathbf{v}\|} X_i$)¹²:

$$\mathbf{Prob}[X > \rho \mathbf{Exp}[X]] < \left(\frac{e^{\rho-1}}{\rho^\rho}\right)^{\mathbf{Exp}[X]},$$

for any $\rho > 1$. In particular, when $\rho \geq 2e$ (where e is the base of the natural logarithm), the latter term is bounded by $2^{-\rho \mathbf{Exp}[X]}$. Recall that we assumed $\|\mathbf{v}\| \leq k$, and thus $\mathbf{Exp}[X] = \mathbf{Exp}[\|\mathbf{v}_{|I_j}\|] \leq k \cdot (m_j - 1)/n$. Thus, for any $t \geq 2e$, we obtain:

$$\mathbf{Prob}\left[\|\mathbf{v}_{|I_j}\| > tk(m_j - 1)/n\right] = \mathbf{Prob}\left[\|\mathbf{v}_{|I_j}\| > \frac{tk(m_j - 1)/n}{\mathbf{Exp}[\|\mathbf{v}_{|I_j}\|]} \cdot \mathbf{Exp}[\|\mathbf{v}_{|I_j}\|]\right],$$

observe that in this case $\rho := \frac{tk(m_j-1)/n}{\mathbf{Exp}[\|\mathbf{v}_{|I_j}\|]} \geq 2e$, due to our assumption on t and the fact that $\frac{k(m_j-1)/n}{\mathbf{Exp}[\|\mathbf{v}_{|I_j}\|]} \geq 1$, and thus the latter term is bounded by:

$$2^{-\frac{tk(m_j-1)/n}{\mathbf{Exp}[\|\mathbf{v}_{|I_j}\|]} \mathbf{Exp}[\|\mathbf{v}_{|I_j}\|]} = 2^{-tk(m_j-1)/n},$$

as asserted. □

E Approximating the Minimum Spanning Tree

Given a set system (X, Σ) as in Section 4.2, our goal is to approximate the spanning tree with minimum conflicts, using the metric embedding approach of Har-Peled and Indyk [16]. The original settings studied

¹²We note that in the original formulation in [31], one needs to have a set of *independent* random variables $\hat{X}_i, i \in \{1, \dots, \|\mathbf{v}\|\}$ with $\hat{X} = \sum_{i=1}^{\|\mathbf{v}\|} \hat{X}_i$, such that $\mathbf{Exp}[X] \leq \mathbf{Exp}[\hat{X}]$. In the scenario of our problem \hat{X}_i is taken to be a Bernulli indicator random variable, which takes value one with probability $(m_j - 1)/n$, in which case $\mathbf{Exp}[X] = \mathbf{Exp}[\hat{X}] = \|\mathbf{v}\| \cdot (m_j - 1)/n$.

in [16] are dual set systems of halfspaces and points (as in [37]), however, most of the steps in their analysis can be applied in our case as well, and therefore we mainly emphasize the modifications required by our analysis.

Specifically, we proceed as follows. For any subset $P \subseteq X$, we form a mapping $f_P : \Sigma \rightarrow \mathbb{Z}$, which maps a region (i.e., a set) in Σ to a unique integer ID in the projection $\Sigma|_P$ (see [16] for the existence of such a mapping). Next, we embed the symmetric difference distance between each pair of sets $S, S' \in \Sigma$ to the Hamming space \mathbb{Z}^μ , where $\mu = \text{polylog}\{m\}$, using a collection \mathcal{P} of μ random subsets P_1, \dots, P_μ of X . Here, the Hamming space consists of all ID vectors $f_{\mathcal{P}}(S) = (f_{P_1}(S), \dots, f_{P_\mu}(S))$, for each $S \in \Sigma$. Then, a main property of the analysis in [16] is to have a low distortion between the symmetric difference distance of pairs of sets S, S' and the Hamming distance between their two corresponding vectors $f_{\mathcal{P}}(S), f_{\mathcal{P}}(S')$ (which is the number of coordinates where they disagree).¹³

Once the mapping $f_{\mathcal{P}}(\cdot)$ is computed, for all sets $S \in \Sigma$, one can use existing machinery to support dynamic approximate nearest-neighbor in order to compute $(1 + \eta)$ -approximation for the spanning tree in time $O(m^{1+1/(1+\eta)})$ (this is described in [16] in detail, see also [12, 22, 25]).

Thus a main task is to compute $f_{\mathcal{P}}(\cdot)$ efficiently. Whereas the computation of $f_{\mathcal{P}}$ in [16] is based on the construction of many faces in an arrangement of lines (for planar settings) and an “intersection-searching” machinery (for settings in higher dimensions), we need to resort to a *range-counting* machinery (see, e.g., [2]). In this case one can obtain a *canonical representation* for the subset of points stored in each $S \in \Sigma$, and then produce the IDs based on the canonical sets (rather than the explicit sets). Skipping the technical (and straightforward) details, we conclude that the running time for computing $f_{\mathcal{P}}(\cdot)$ is dominated, up to a polylogarithmic factor, by the task of computing, for each $S \in \Sigma$, the number of points that it contains from P_1, \dots, P_μ . By existing range searching machinery (see, e.g., [2]), it follows that this time is (i) $O^*(n^{d/(d+1)}m^{d/(d+1)} + n + m)$, if Σ is a set of halfspaces or parallel slabs in d -space, and (ii) $O^*(n^{(d+1)/(d+2)}m^{(d+1)/(d+2)} + n + m)$, if Σ is a set of balls in d -space. Thus from the above considerations we obtain the bounds stated in Corollary 4.6.

¹³We note that this part of the analysis in [16] is sufficiently general to include our setting as well.